

Cambria Cluster User Manual

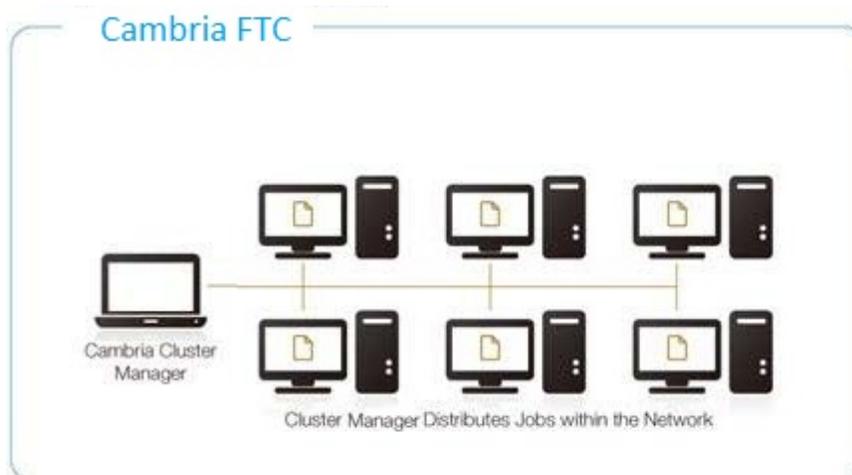
Cambria Operation - User Guide

I. Introduction	3
1. System and Network Requirements	3
a. Minimum System Requirements – Cambria Cluster Manager	4
b. Minimum System Requirements – Encoding Machine	4
c. Network Requirements	4
1. Operational Paradigm	4
a. Job Distribution Within a Cluster	4
I. Installation and Setup	5
1. Controlling Encoding machines on the Cluster	5
2. Disconnecting Encoding Machines from the Cluster	6
3. Setting Encoding Machine Priority	7
4. Choosing the Number of Slots	8
5. Working with Dynamic Slot	8
6. Working with Job Tags	9
a. Working with Job Tags	9
b. Sample Tagging Scenario	10
c. Tagging Machines in Cambria Cluster Manager	11
d. Tagging Jobs in the Preset Editor and Setting Priority	12
1. Advanced Operations	12
a. CPU Baseline Tag	12
b. CPU Baseline Adjust Percentage	13
c. Assigning Jobs from Encoding Machines to the Cluster	14

I. Introduction

Cambria Cluster is an application that allows management of other Cambria FTC's with a single manager client. Users can encode many jobs simultaneously by leveraging all available machines with Cambria FTC installed. Job distribution and synchronization are handled automatically by the manager.

The Cluster is also highly expandable and client machines can be added or removed at any time without affecting other existing client machines. Here is a diagram of Cambria Cluster Manager at the head of a group of File Convert machines.



A. System and Network Requirements

As shown in the figure, there are two types of machines in a cluster; the Cluster Manager, which runs Cambria Cluster Manager, and encoding machines, which run the standard version of Cambria FTC.

Though the Cambria Cluster manager software can perform the same encoding tasks as Cambria FTC, for most efficient operation, Capella recommends that the Cluster Manager **not** be used for encoding jobs with more than five FTC's. If encodes are to be run on the Cluster Manager, it should match the encoding machine specs shown below, but should use only one encoding slot maximum. More on this below.

1. Minimum System Requirements – Cambria Cluster Manager

Operating System: Windows 10 (64 bit), Windows Server 2012 R2, Windows Server 2016

*** File Convert and Cluster installation requires addition system files included in recent Windows Update.

*** If you encounter issues during installation, please conduct "Windows Update", reboot, and try again.

Processor: Intel 2.8GHz (Quad-Core) or faster

Memory: 8GB or more (4K/8K encoding requires more memory 32GB suggested)

Video Card: Supports Direct3D acceleration

2. Minimum System Requirements – Encoding Machine

Operating System: Windows 10 (64 bit), Windows Server 2012 R2, Windows Server 2016

*** File Convert and Cluster installation requires addition system files included in recent Windows Update.

*** If you encounter issues during installation, please conduct "Windows Update", reboot, and try again.

Processor: Intel 2.8GHz (Quad-Core) or faster

Memory: 8GB or more (4K/8K encoding requires more memory 32GB suggested)

Video Card: Supports Direct3D acceleration

3. Network Requirements

- a. All Client machines must be in same subnet as the Cluster Management machine.
- b. The cluster manager and all encoding machines must be connected on the same wired LAN. Wireless connections will severely degrade cluster performance.
- c. IPV6 is not supported. Please use IPV4 networks.
- d. Network shares/drives are required. They should be accessible with proper permissions by Client machines.

4. Operational Paradigm

Most clusters can utilize watch folders and API via the Cluster Manager which are accessed by the various system users. These jobs are distributed among the encoding machines according to rules discussed below.

However, each encoding machine is a fully functional version of Cambria FTC that can run jobs via the UI, watch folders, or API. Encoding machines can be easily disconnected from the cluster to maximize efficiency of jobs running on that station, and each encoding machine can send jobs to the cluster.

1. Job Distribution Within a Cluster

You'll learn how to set these controls below, but let's take a high level look at how Cluster Manager distributes jobs within the cluster. As with all Cambria encodes, you can set a priority for each job that will control when that job is processed in relation to other jobs in the system.

Status	Machine name	Machine priority	CPUs	Slots	Use Dynamic Slot	Slots in use	CPU Usage (%)	Succeeded/Total	Tags
Online	Z840-HP	1 (Highest)	40	3	<input type="checkbox"/>	3	100	100% (2/2)	FastMachine
Online	Z600	5 (Normal)	16	2	<input type="checkbox"/>	0	0	0% (0/64)	HEVC

Figure 1. These options control job distributions within the cluster.

Once it's time for the job to be assigned to an encoding machines, several concepts control which machine is assigned.

Machine priority – As you can see in the figure above, you can assign a priority to each system in the cluster. Jobs are first distributed to machines with the highest priority.

Slots – For each encoding machine, you'll set the number of slots for concurrent encoding jobs. If the machine with the highest priority has an open slot, the job will be assigned to that machine. If several machines have the same priority and same number of open slots, the job will be assigned arbitrarily.

Dynamic slots – If you enable Dynamic Slots, the Cluster Manager will use the machine's average CPU percentage over the previous 20 secs to determine if it can accept more jobs to run concurrently. Once the average CPU falls below 90%, Cluster Manager will assign another job to that machine until it reaches the maximum number of slots for that encoding machine. As you will learn below, there are several adjustments that you can make to dynamic slots to balance cluster performance and utilization.

Tags – On the extreme right of the figure, you see a column for Tags and tags in the two encoding machines shown. Tags allow you to match jobs with particular encoding machines, whether for performance or to access features not available on all encoding machines.

II. Installation and Setup

Install Cambria Cluster Manager on the cluster controller, and Cambria FTC on all individual encoding machines. Please refer to the release notes for more information regarding the specifications and credentials of your version of Cambria Cluster.

A. Controlling Encoding machines on the Cluster

Once you install the cluster, encoding machines that meet the network requirements presented above should appear in the Machines Tab as shown in Figure 2. Initially, encoding machines will be in standby mode, so you can test its configuration and generally get it up and running before being assigned jobs from the cluster.

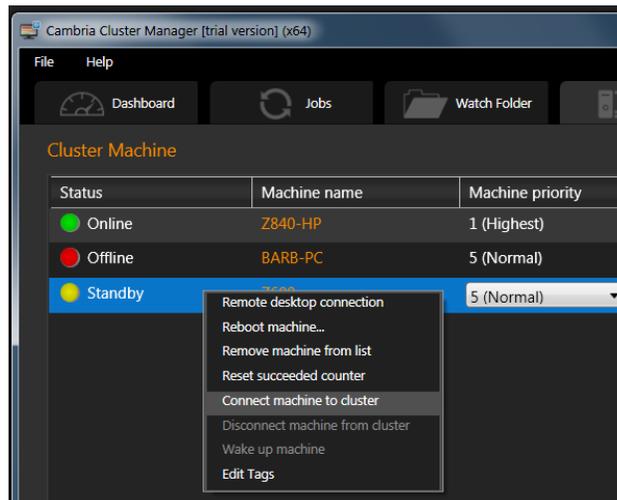


Figure 2. Connecting and disconnecting machines from the cluster.

Computers in the cluster will appear with an indicator to its left as shown in Figure 2. Here's what the color codes mean.

- **Online**: Machine is reachable, Cluster will assign jobs to this machine.
- **Offline**: Machine is not reachable. Cluster will not assign jobs to this machine.
- **Standby**: Machine is reachable, but not connected to Cluster. Cluster will not assign jobs to this machine.
- **Cold Standby**: Machine is not reachable, and not connected to Cluster. Cluster will not assign jobs even if machine is reachable in the future.

1. To connect a computer to the cluster, and make it accessible by the cluster, right click the computer and choose Connect Machine to Cluster. The yellow indicator should change to green.

B. Disconnecting Encoding Machines from the Cluster

By default, encoding machines are placed in standby mode until connected as shown above. To disconnect the encoding machine from the cluster, take the following steps.

1. Run Cambria Manager on the encoding machine.
2. Choose File > Options to open the Options dialog on the encoding machine.
3. In the Options dialog, deselect the Broadcast to Cambria Cluster checkbox in the Cluster settings area (Figure 3). Once you do this, the machine will show up as offline with a red indicator in the Cluster Manager on the Cluster Manager.

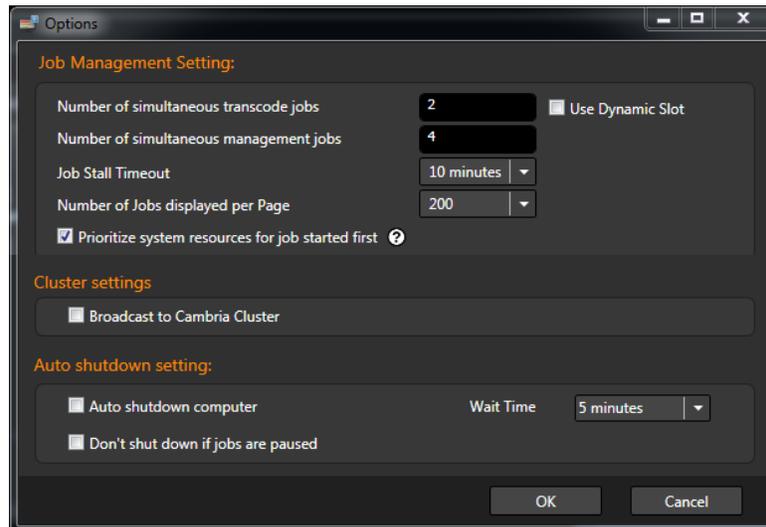


Figure 3. The Options dialog on the encoding machine.

C. Setting Encoding Machine Priority

As discussed above, machine priority controls the order in which jobs are sent to an encoding machine in the cluster, with jobs sent to the highest priority encoding machine with an open slot. Typically, you'll want to give the highest priorities to your most capable encoding machines.

To set priority for each encoding machine, take the following steps:

1. Run Cambria Cluster Manager.
2. Choose the Machines tab.
3. Click the drop-down list for the encoding station and choose the desired machine priority (Figure 4). Note that lower numbers designate higher priority levels.

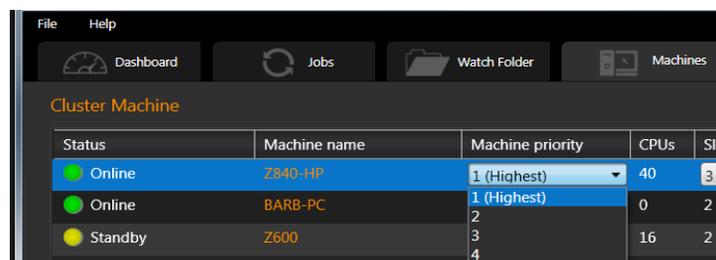


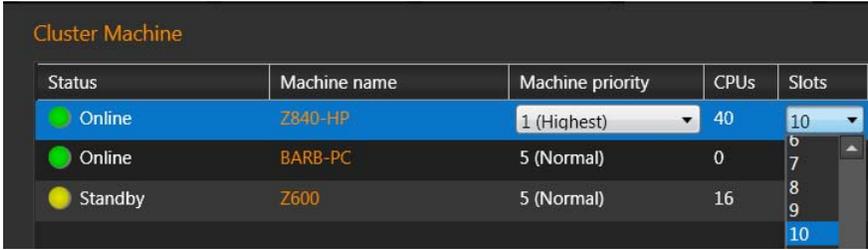
Figure 4. Selecting Machine priority.

D. Choosing the Number of Slots

Slots control the number of jobs that can be assigned to each encoding machine. Note that all jobs are not created equal, so each job won't consume the same CPU cycles as another. A single job encoding a long 4K file to nine adaptive bitrate levels using the HEVC codec might consume all the available CPU cycles in a 40-core computer for several hours, while encoding a short 720p file to MPEG-2 format might consume a fraction of the CPU cycles for just a few moments. So, assigning jobs based upon slots is inherently inefficient, which makes Dynamic Slots, discussed next, a better option for most clusters.

If you decide against using Dynamic Slots, however, consider the following guidelines when setting the available slots.

- The more CPUs available on the encoding machine, the more slots you should assign.
- When encoding on your cluster manager, never assign more than one slot to that computer.



Status	Machine name	Machine priority	CPUs	Slots
Online	Z840-HP	1 (Highest)	40	10
Online	BARB-PC	5 (Normal)	0	7
Standby	Z600	5 (Normal)	16	8
				9
				10

Figure 5. Choosing the number of slots.

To set the number of slots for each encoding machine, take the following steps:

1. Run Cambria Cluster Manager.
2. Choose the Machines tab.
3. Click the drop-down list for the encoding station and choose the desired machine priority (Figure 5).

E. Working with Dynamic Slot

When you enable Dynamic Slot, the Cluster Manager assigns jobs based upon available CPU in addition to the number of available slots. Specifically, when you check "Use Dynamic Slot," Cluster Manager will check the encoding machine's average CPU percentage over the previous 20 seconds. If the average CPU percentage exceeds 90%, Cluster Manager won't send that encoding machine

additional jobs. If the average CPU falls below 90% for that period, Cluster Manager will send additional jobs to that encoding machine until the number of running jobs equals the selected slot number.

In an encoding environment where there are always queued jobs awaiting available encoding resources, Dynamic Slot allows Cluster Manager to adjust the number of simultaneous transcode jobs to sustain a CPU utilization of 90% to 100%. This maximizes overall cluster CPU utilization, while ensuring that you don't compromise the speed of any individual job by assigning too many jobs to any single machine. Dynamic Slot works well with varying types of jobs on clusters with disparate encoding machine specifications.



Status	Machine name	Machine priority	CPUs	Slots	Use Dynamic Slot
Online	Z840-HP	1 (Highest)	40	15	<input checked="" type="checkbox"/>

Figure 6. Enabling Dynamic Slot.

To enable Dynamic Slot, take the following steps.

1. Run Cambria Cluster Manager.
2. Choose the Machines tab.
3. Click the Use Dynamic Slot checkbox for the target encoding station.
4. Consider increasing the number of slots to ensure that CPU utilization can reach 90% or higher during peak encoding periods. Again, Cluster Manager won't send additional jobs to an encoding machine if it has no open slots even if CPU utilization is under 90%. If the number of slots is set too low, a powerful computer may never be assigned sufficient jobs to maximize CPU utilization.

F. Working with Job Tags

Job tags allow Cluster users to match specific jobs to specific encoding machines, whether based upon performance, available features, or other criteria. Operationally, you tag different computers in the Machines tab in Cambria Cluster Manager, and then tag individual jobs or watch folders using the Preset Editor. Let's discuss how job tags work before demonstrating how to apply these tags.

1. Working with Job Tags

Here are the rules for creating and working with job tags.

- a. Jobs are configured with text tags, with optional connectors. The “&” connector means that **both** tags must be satisfied, while the “|” connector means **either** tag must be satisfied.
- b. Encoding machines are also configured with text tags separated by commas or semi-colons.
- c. Tags are used as a matching or non-matching criteria when jobs are accepted by the Cluster Manager.
- d. Tags can be used as an inclusive or exclusive match.
- e. Tags should not be case sensitive.

2. Sample Tagging Scenario

This scenario demonstrates how these rules operate.

- a. Machine Tags. This cluster has three encoding machines tagged as follows.

PC1 - Fast, RAID, Cineform, Urgent

PC2 - HEVC, Long

- b. Here are the example jobs, submitted to the Cluster Manager in this order with the following tags. Example jobs

Job1 – Fast (assign to an encoding machine with a Fast tag)

Job2 - !Fast (assign to any encoding machine without a Fast tag)

Job3 – Cineform (assign to an encoding machine with the Cineform tag) **Job4** -

Fast&RAID (assign to an encoding machine with a Fast **and** RAID tag) **Job5** -

Fast|Long (assign to an encoding machine with a Fast **or** Long tag) **Job6** – Flash
(assign to an encoding machine with the Flash tag)

Job7 - ... no tag ... (assign to any encoding machine)

- c. Here are the job assignments if all encoding machines have the same priority and two open transcoding slots.

Job1 - PC1 (PC1 has the Fast tag and an open slot. This is the first job assigned to PC1)

Job2 - PC2 (can go to either PC2 or PC3 since neither have the Fast tag. Assigned to PC2 because next in line)

Job3 - PC1 (PC1 has the Cineform tag and an open slot. This is the second job assigned to PC1, which now can't accept additional jobs until finishing a job)

Job4 - Still in queue (tried PC1 because it's the only machine with Fast and Raid, but there is no free slot)

Job5 - PC2 (Can go to machine with either Fast or Long; PC2 has Long plus an open slot).

Job6 - PC3 (can go to any machine; PC3 has the only slot).

3. Tagging Machines in Cambria Cluster Manager

To tag machines, take the following steps.

1. Run Cambria Cluster Manager.
2. Choose the Machines tab.
3. Right click the target encoding machine, and choose Edit Tags (Figure 7).

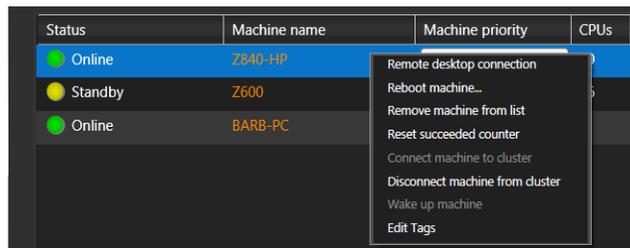


Figure 7. Opening the Edit Tags dialog.

4. In the Edit Tags dialog, type the desired tags, separated by a comma or semi-colon (Figure 8).

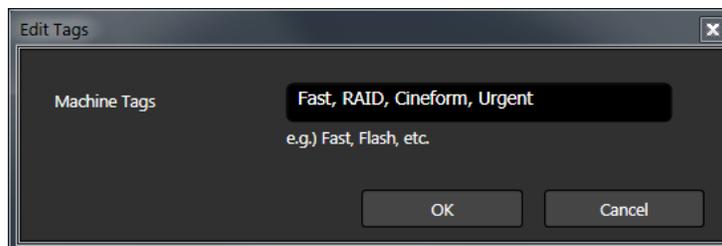


Figure 8. Adding tags to an encoding machine.

- Click OK to close the dialog and repeat as needed (Figure 9).

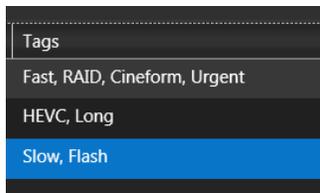


Figure 9. Our machines, tagged as above.

4. Tagging Jobs in the Preset Editor and Setting Priority

Whether you're creating a job in Cambria FTC or in a watch folder, you create tags and set encoding priority in the Preset Editor. To do so, follow these steps.

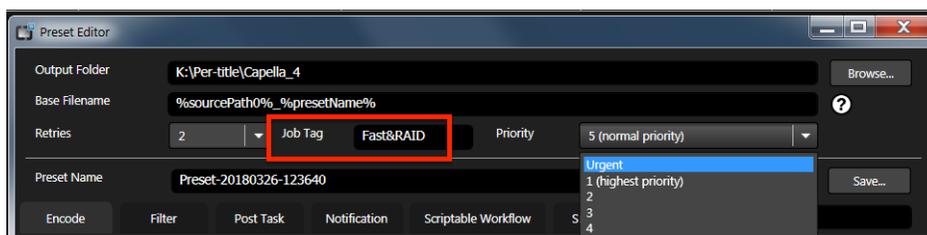


Figure 10. Adding job tags and setting encoding priority in the preset editor.

- To add tags, simply type them into the Job Tag field highlighted in Figure 10, along with any modifiers.
- To set encoding priority, use the drop-down list on the upper right. Note that jobs designated as Urgent are immediately assigned to the highest priority encoding machine compatible with any tags, and all other running jobs are paused on that encoding machine until the urgent job is completed.

III. Advanced Operations

In this section, we'll cover two advanced tags designed to be used when Dynamic Slot is enabled, and how to assign a job from an encoding machine to a cluster.

A. CPU Baseline Tag

Many encoding jobs, particularly those that use 2-pass encoding, use different levels of CPU cycles over the duration of the job. During 2-pass encoding, for example, CPU utilization is very low during the first analysis pass, but very high during the second encoding pass. If Cluster Manager checks CPU utilization during the first pass, it may assign additional jobs that cause an over-distribution of jobs to an encoding machine.

A new job tag can prevent this from occurring. Specifically, you can insert the job tag 'CPUBaseline=X' where X is a user-specified number from 1 and 100. When the job is running, Cluster Manager will use the X value as the CPU utilization value for that Job, and will aggregate the X value for all jobs running on an encoding machine to decide whether to assign any additional jobs to the machine.

For example, assume you inserted CPUBaseline=80 into a job. When this job is running, it will report to Cluster that the Job is using 80% CPU utilization at the very least. If the Job is using more than 80% CPU, the higher value will be reported).

Like the job tags shown above, this tag is inserted in the Preset Editor, with comma or semi-colon separating it from other tags (see Figure 11).

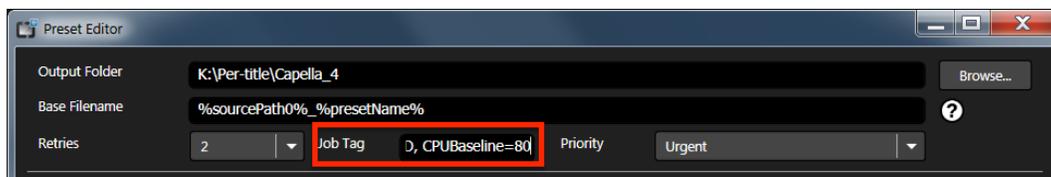


Figure 11. Adding the CPUBaseline tag.

B. CPU Baseline Adjust Percentage

The CPUBaseline tag works well in a homogenous encoding farm where all machines have similar processing power, but what about when they differ? A job might take 80% of CPU on a 4-core computer, but only 20% on a 40-core computer. Without a way to compensate, the CPU Baseline tag could lead to serious under-utilization on more powerful computers.

According, Cluster Manager can deploy a new special Machine Tag to dynamically adjust CPUBaseline for the processing power of the machine upon which the job is running. Specifically, you can set the machine tag 'CPUBaselineAdjPct=X' through the Cluster Manager for any Machine in the machines list, with X a number from 0 to 100.

This value scales the CPU Baseline value for each machine. As an example, in the cluster shown above, we might assign a CPU Baseline Adjust Percentage of 25 to the 40-core HP Z840. If Cluster Manager assigned a CPUBaseline=80 job to the 40-core HP, Cluster Manager would multiply the Baseline Percentage of 80 by .25 and set CPU utilization at 20% for the purposes of assigning additional jobs to that computer. Of course, if actual CPU utilization was higher than the aggregate value of all adjusted job baselines, Cluster Manager would use the actual value.

This is a machine-specific vale that you set as shown in Tagging Machines in Cambria Cluster Manager above. Follow the instructions there to open the Edit Tags

dialog, and type the desired CPUBaselineAdjPct in the dialog. Click OK to close the dialog.

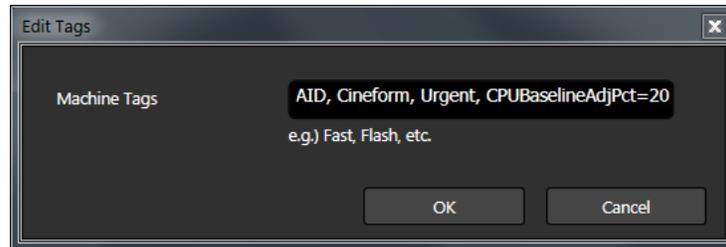


Figure 12. Adding the CPUBaselineAdjPct tag to compensate for a faster computer.

C. Assigning Jobs from Encoding Machines to the Cluster

As mentioned above, jobs can be launched from each encoding machine for local processing. At times, however, operators may wish to feed jobs to the cluster to preserve local CPU resources or accelerate performance. Follow these steps to do so.

1. In Cambria FTC's encoding tab, create the job as normal. Then click the downward arrow on the Queue All Jobs button and choose either Queue All Jobs to Remote Machine or Queue Selected Jobs to Remote Machine (Figure 13).

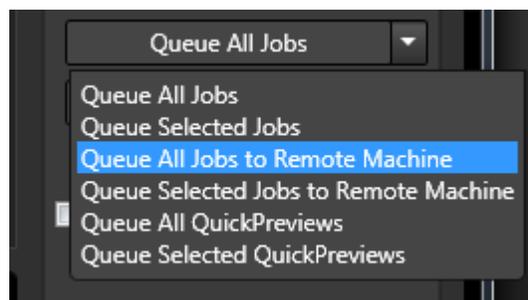


Figure 13. Sending a job to the cluster.

2. Cambria opens the Select Remote Machine Dialog. In that dialog, type the local address for the Cluster Manager (Figure 14). If you don't know the address, open a browser and navigate to www.whatismyip.com.

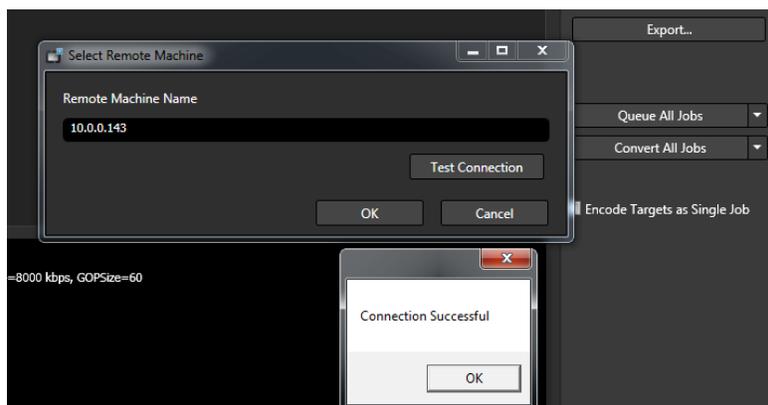


Figure 14. Entering the remote machine address and testing the connection.

3. Click the Test Connection button to test the connection. If successful, click OK to close the dialog which will run the job. If not, enter the correct IP address and try again.